



DGI2 Fachgebiet 5-2 – Accounting

# Installation des DGAS-Clients für das D-Grid Accounting

Jan Wiebelitz (wiebelitz@rvs.uni-hannover.de)

Michael Brenner (brenner@rvs.uni-hannover.de)

29.02.2009

Version 0.8

## Inhaltsverzeichnis

1	INSTALLATION DES DGAS-CLIENT .....	3
2	KONFIGURATION VON DGAS (CLIENT).....	3
3	ANPASSUNGEN AN DGAS .....	3
3.1	grid-mapfile.dgas_special .....	3
3.2	groupmapfile.dgas .....	5
3.3	Erweiterung von glite-urcollector.pl .....	5
4	DIENSTE STARTEN.....	8
5	KONFIGURATIONSDATEIEN .....	9
5.1	dgas_atmClient.conf .....	9
5.2	dgas_ce_getAcctLogd.conf.....	10
5.3	dgas_ce_pushd.conf .....	11
5.4	dgas_gianduia.conf.....	12
5.5	dgas_VOUserHLR.conf .....	16

## 1 Installation des DGAS-Client

Diese Installationsanleitung beschreibt die Installation des DGAS-Clients auf Basis von gLite 3.1 und Scientific Linux CERN (SLC) 4.6. Die Installationsbeispiele beziehen sich auf der Installation am RRZN in Hannover, d.h. dass die Rechnernamen in den Beispielen gegen lokale Rechnernamen ausgetauscht werden müssen.

Die Installation des DGAS-Clients basiert auf der Installation eines LCG-CE, da einige Paketabhängigkeiten bestehen. Auf dem LCG-CE werden folgende Pakete installiert:

- glite-dgas-common
- glite-dgas-hlr-clients
- lcg-dgas-tools

```
# yum install glite-dgas-common.i386 glite-dgas-hlr-clients.i386 lcg-dgas-tools.i386
```

## 2 Konfiguration von DGAS (Client)

Die Konfigurationsdateien befinden sich im Verzeichnis */opt/glite/etc*

```
# cd /opt/glite/etc
```

Folgende Dateien müssen angepasst werden (siehe weiter unten):

- dgas\_atmClient.conf
- dgas\_ce\_getAcctLogd.conf
- dgas\_ce\_pushd.conf
- dgas\_gianduia.conf
- dgas\_VOUserHLR.conf

Das Verzeichnis mit den Torque Accounting Informationen muss auf dem DGAS-Client verfügbar gemacht werden, Beispiel:

```
# mkdir -P /var/spool/torque/server_priv/accounting
# mount -t nfs torqued1.d-grid.uni-
hannover.de:/var/spool/torque/server_priv/accounting
/var/spool/torque/server_priv/accounting
```

## 3 Anpassungen an DGAS

Damit die richtigen Nutzerinformationen dem Accounting-Record zugeordnet werden, unabhängig über welche Middleware ein Job abgegeben wurde, ist für das Mapping des Accounts auf den zugehörigen Distinguished Name (DN) eine Datei namens *grid-mapfile.dgas\_special* vorhanden.

Da im D-Grid jeder Nutzer auf einen eigenen Account abgebildet wird, kann diese statische Zuordnung genutzt werden, um einem Nutzer die Accounting-Informationen zu einem Job zuzuordnen.

### 3.1 grid-mapfile.dgas\_special

Die Datei *grid-mapfile.dgas\_special* befindet sich im Verzeichnis */etc/grid-security* und dient dazu, den Account eines Nutzers dessen DN zuzuordnen. Da in der DGAS-Datenbank, dem Home Location Register (HLR), jeder DN eindeutig ist, wird der DN um die VO des Nutzers erweitert.

```
"/O=GermanGrid/OU=UniHannover/CN=Jan Wiebelitz <dgtest>" dt0024
```

Das D-Grid Grid-Mapfile muss aus Jülich bezogen werden. Dazu existiert ein entsprechender Cron-Job, der diesen Vorgang automatisiert.

*/etc/cron.d/getgridmap:*

```
PATH=/sbin:/bin:/usr/sbin:/usr/bin
2 2,8,14,20 * * * root /usr/sbin/getgridmap >> /var/log/getgridmap 2>&1
```

Der Vorgang wird durch das Skript `getgridmap` (gleicher Name, aufgepasst) ausgeführt.

*/usr/sbin/getgridmap*

```
#!/bin/bash

DATE=`date +%Y%m%d%s`

cd /etc/grid-security/preprocess
cp grid-mapfile.dgas_special grid-mapfile.dgas_special.$DATE

echo "-----" `date` " dgridmap"
# -output-g-old setzt vor den Ausgabedateinamen "./"
/usr/sbin/dgridmap -pre "" -cert-path /etc/grid-security/lcgced1/ -output-g-old grid-mapfile.dgas_new

echo "insertGroup.pl"
/usr/sbin/insertGroup.pl grid-mapfile.dgas_new ../groupmapfile.dgas grid-mapfile.dgas_dnmap > grid-
mapfile.dgas_special

# hier nur den erforderlichen localuser anfüegen
#echo "/C=DE/O=GermanGrid/OU=UniHannover/CN=localuser <localuser>" localgroup' >> grid-
mapfile.dgas_special
#echo "/O=GermanGrid/OU=Uni-Dortmund/CN=localuser <localuser>" localgroup' >> grid-mapfile.dgas_special
#echo "/C=DE/O=GridGermany/OU=Forschungszentrum Juelich GmbH/OU=ZAM/CN=localuser <localuser>" localgroup'
>> grid-mapfile.dgas_special

cp grid-mapfile.dgas_special ..
```

Sofern erforderlich, kann im `getgridmap`-Skript ein Local-User an das Gridmap-File angefügt werden (siehe Kommentar).

Die Aufbereitung der Datei findet im anzulegenden Verzeichnis `/etc/grid-security/preprocess` statt. Dabei fügt das Skript `/usr/sbin/insertGroup.pl` die VO-Namen in die DN der Benutzer ein.

Das `grid-mapfile.dgas_special` sollte dann wie folgt aussehen:

```
...
"/O=GermanGrid/OU=UniHannover/CN=Christian Grimm <gdigrid>" gd0005
"/O=GermanGrid/OU=UniHannover/CN=Christian Grimm <kerndgrid>" kg0012
"/O=GermanGrid/OU=UniHannover/CN=Christopher Kunz <education>" ed0008
"/O=GermanGrid/OU=UniHannover/CN=Christopher Kunz <gdigrid>" gd0006
"/O=GermanGrid/OU=UniHannover/CN=Denis Goehr <dgtest>" dt0008
"/O=GermanGrid/OU=UniHannover/CN=Denis Goehr <gdigrid>" gd0007
"/O=GermanGrid/OU=UniHannover/CN=Denis Goehr <kerndgrid>" kg0009
"/O=GermanGrid/OU=UniHannover/CN=Gian Luca Volpato <education>" ed0018
"/O=GermanGrid/OU=UniHannover/CN=Gian Luca Volpato <kerndgrid>" kg0015
"/O=GermanGrid/OU=UniHannover/CN=Harald Schwier <education>" ed0020
"/O=GermanGrid/OU=UniHannover/CN=Harald Schwier <kerndgrid>" kg0011
"/O=GermanGrid/OU=UniHannover/CN=Jan Wiebelitz <dgtest>" dt0024
```

...

### 3.2 groupmapfile.dgas

Die Datei *groupmapfile.dgas* befindet sich im Verzeichnis */etc/grid-security* und dient zur Umsetzung der VO-Kürzel in VO-Namen

```
as astrogrid
ad astrogrid
bw bwgrid
c3 c3grid
dt dgtest
ed education
fi fingrid
gd gdigrid
hp hepcg
in ingrid
kg kerndgrid
ls lifescience
md medigrid
pa partnergrid
pr progrid
rv rvs
tx textgrid
vi viola
ws wisent
```

Neue VOs müssen in dieser Datei eingepflegt werden.

### 3.3 Erweiterung von *glite-urcollector.pl*

In dem Perlskript */opt/glite/sbin/glite-urcollector.pl* ist die Funktion *writeGianduiottoFile* an die Gegebenheiten von D-Grid angepasst worden:

```
## ----- write file in gianduiottiBox per pushd ----- ##
sub writeGianduiottoFile {

    my $filename = $lrmsType."_".$_[0]."_".$_[1];
    # unique filename: <lrmsType>_<lrmsJobID>_<lrmsEventTimestamp>
    my $header = $_[2];
    my $acctlog = $_[3];

### JAN ab hier
    my $gridmapfile="/etc/grid-security/grid-mapfile.dgas_special";
    my $groupmapfile="/etc/grid-security/groupmapfile.dgas";
    my $trash = "";
    my $localuser = "";
    my $localgroup = "";
```

```

my $user = "";
my $group = "";
my $DN = "";
my $vo = "";
my $vokrzl = "";
my $acctline = "";
my $acctdata = "";

$acctline = $acctlog;
$acctline =~ s/\;/ /;

### JAN ab hier
#
# wir reiten ueber die accounting-saetze aus dem pbs und ersetzen <user> und <group>
# entsprechend der regeln:
#
# 1. praefix bei einem [aaaannnn]-format wird immer entfernt -> uhad0006 wird zu
ad0006
#
# 2. fuer den user wird ein DN im gridmapfile gesucht
#
# 2a. DN gefunden -> vo-name anhand vo-kuerzel des users ermitteln (groupmapfile)
# 2b. kein DN gefunden -> <user> ist 'localuser', <group> ist 'localgroup'
#
foreach $acctdata ( split /\s/, $acctline ) {
    if ( $acctdata =~ m/user/ ) {
        ( $trash, $localuser ) = split /=/, $acctdata;

        $user = $localuser;

        # 1.
        if ( $localuser =~ /\w\w\w\w\d\d\d\d/ ) {
            $user =~ s/\w\w(\w+)/$1/;
        }

        # 2.
        $DN = ( split "/", `grep $user $gridmapfile` )[1];

        # 2b.
        if ( $DN eq "" ) {
            $user = "localuser";
            $group = "localgroup";
            $DN = ( split "/", `grep localuser $gridmapfile` )[1];
        }
    }
}

```

```

        else    # 2a.
        {
            $vokrzl = substr( $user, 0, 2 ) . " ";
            $group = ( split / /, `grep "$vokrzl" $groupmapfile` )[1];
            chomp $group;
        }

        }
        if ( $acctdata =~ m/group/ ) {
            ( $trash, $localgroup ) = split /=/, $acctdata;
        }
    }

$acctlog =~ s/group=$localgroup/group=$group/;
if ( $user eq "localuser" ) {
    $acctlog =~ s/user=$localuser/user=localuser/;
}
else {
    $acctlog =~ s/user=$localuser/user=$user/;
}
### JAN bis hier

    open(OUT, "> $gianduiottiBox/$filename") || return 1;

    print OUT "$header";
### Jan ab hier
    print OUT "USER_DN=$DN\n";
    print OUT "USER_FQAN=/$group/Role=NULL/Capability=NULL\n";
### Jan bis hier
    print OUT "ACCTLOG:$acctlog";

    close (OUT);

    return 0;
}

```

Diese Funktion muss in der installierten Version von `/opt/glite/sbin/glite-urcollector.pl` ausgetauscht werden.

### 3.4 SPEC Benchmark

Um die Ressourcennutzung zwischen unterschiedlichen Ressourcen vergleichbar zu gestalten, ist in DGAS die Nutzung der SPECint2000 und SPECfp2000 Benchmarks vorgesehen, da die Zahlen für diesen Benchmark nicht mehr auf der Webseite von [spec.org](http://www.spec.org) veröffentlicht werden, wurden für die Nutzung in D-Grid die Benchmarks SPECint2006 und SPWCfp2006 gewählt. Für eine Vielzahl von Systemen sind diese Benchmarks veröffentlicht (<http://www.spec.org/cpu2006/results/cpu2006.html>). Damit für die eigenen Ressourcen die passenden Benchmarks in das Home Location Register eingetragen werden,

muss die Datei `dgas_cpu_bench.conf` angelegt werden und die passenden Benchmarks eingetragen werden. Beispiel:

```
# SPECint_base2006
GlueHostBenchmarkSF00:17.1
GlueHostBenchmarkSI00:17.3
```

In der Konfigurationsdatei `/opt/glite/etc/dgas_gianduia.conf` muss diese Datei referenziert werden (siehe unten).

#### 4 Dienste starten

Vor dem Starten der Dienste muss die Ressource (der CE) noch als Ressource in der DGAS-Datenbank eingetragen werden. Hierzu bitte den DN aus dem Server-Zertifikat an [accounting@rvs.uni-hannover.de](mailto:accounting@rvs.uni-hannover.de) mailen.

Die DGAS-Startskripte in das Verzeichnis `/etc/init.d` kopieren:

```
# cp /opt/glite/sbin/glite-dgas-urcollector /etc/init.d
# cp /opt/glite/sbin/glite-dgas-pushd /etc/init.d
```

und mit `chkconfig` aktivieren, so dass die DGAS-Dienste beim Neustart des CEs neu gestartet werden:

```
# chkconfig glite-dgas-urcollector on
# chkconfig glite-dgas-pushd on
```

Um Ausreißer zu finden bitte vor dem Start der Dienste eine Auswahl an Accounting-Logs an [accounting@rvs.uni-hannover.de](mailto:accounting@rvs.uni-hannover.de) schicken, damit evtl. noch Anpassungen an der Funktion `writeGianduiottoFile` vorgenommen werden können.

## 5 Konfigurationsdateien

### 5.1 dgas\_atmClient.conf

```
# This is the configuration file for the ATMClient API, part of
# gLite DGAS, DataGrid Accounting System.
# Author: A.Guarise -- andrea.guarise@to.infn.it
# Author: R.M.Piro -- piro@to.infn.it
# Author: G.Patania -- patania@to.infn.it

#
#You should modify the following parameters according to your needs.
#

# next parameters specify the full contact string for the Resource
# PA and HLR.

#This configuration file is reed by the DGAS glite_dgas_atmClient api
#And the info the following parameters are passed to the HLR of the user
#that has to be debited for a running job. These info are then used by that
#HLR to correctly process the economic transaction.

# resource PA: in the form host:port:X509_certSubject
# res_acct_PA_id = "hostname:portnumber:X509CertSubject"
# resource HLR: in the form host:port:X509_certSubject
# res_acct_bank_id = "hostname:portnumber:X509CertSubject"
res_acct_bank_id = "ref-dgas.d-grid.uni-
hannover.de:56568:/C=DE/O=GermanGrid/OU=UniHannover/CN=ref-dgas.d-grid.uni-
hannover.de"

# specify wheter do we want economic accounting or not.
economicAccounting = "no"
```

## 5.2 dgas\_ce\_getAcctLogd.conf

Standardkonfiguration, es wurde nichts geändert!

```
listenerActive = "false"
aclFile = "/opt/glite/etc/getAcctLogd.acl"
listeningPort = "56565"
logfile = "/opt/glite/var/log/getAcctLogd.log"
lockFile = "/opt/glite/var/getAcctLogd.lock"
# if outputDir is specified instead of outputFile, the file name
# is received from the remote host.
outputFile = "/opt/glite/var/"
outputDir = "/opt/glite/var/"
```

### 5.3 dgas\_ce\_pushd.conf

Bitte bei *siteName* das Kürzel der eigenen Institution ohne Leerzeichen eintragen

```
# This is the configuration file for the dgas_ce_pushd daemon, part of
# gLite DGAS, DataGrid Accounting System.
# Author: A.Guarise -- andrea.guarise@to.infn.it
# Author: R.M.Piro -- piro@to.infn.it
# Author: G.Patania -- patania@to.infn.it

#
#You should modify the following parameters according to your needs.
#

#The following parameters are needed by the glite_dgas_ce_pushd daemon.
#The daemon is recovers the UR of the job and asynchronously communicates
#them to the User HLR vie the glite_dgas_atmClient api.
#
dgasURDir = "/opt/glite/var/dgasURBox/"
dgasErrDir = "/opt/glite/var/dgasURBox/ERR/"
qDepth = "4"
qMult = "3"
lockFileName = "/opt/glite/var/dgas_ce_pushd.lock"
logFileName = "/opt/glite/var/log/dgas_ce_pushd.log"
mainPollInterval = "10"
queuePollInterval = "50"

#This parameter points to a secondary configuration file that contains a
#list of default User HLRs for the single VOs. These default User HLR contact
#strings are used if no User HLR is specified with the job's JDL expression.
#In case no VO-specific default User HLR is specified the defaultUserHLR can
#be used.
#
#defaultVOUserHLRlist = "/opt/glite/etc/dgas_VOUserHLR.conf"

#This is the User HLR where, by default, user records are sent.
#It is used only if no User HLR is specified with the job's JDL expression and
#if there is no default specifically for the user's VO (see parameter
#defaultVOUserHLRlist).
```

```

#Change it according to your needs, or comment it out if you do not
#need a default userHLR.
#
#defaultUserHLR = "<userHLR_hostname>:<port>:<host_cert_subject>"defaultUserHLR =
"ref-dgas.d-grid.uni-hannover.de:56568:/C=DE/O=GermanGrid/OU=UniHannover/CN=ref-
dgas.d-grid.uni-hannover.de"

#This specifies if UR shall be archived only in the resource HLR ("yes")
#or in the User HLR as well ("no")
#
forceLocalOnly = "no"

#This is the user used by the daemon to run the unprivileged part
#of the script. it MUST not be 'root'. If it is empty, teh pool account
#of the user that runned the job will be used.
#
gridUser = "dgas"

#siteName parameter can be used to assign to the UsageRecord a tag identifying
# the site. Example: "INFN-TORINO".
siteName = "Example-RZ"

# This defines whether to use the name of the host on which the sensors run
# as the host name of the CE (in case of local jobs), default is "no":
localHostNameAsCEHostName = "no"

```

## 5.4 dgas\_gianduia.conf

Bitte bei *siteName* das Kürzel der eigenen Institution ohne Leerzeichen eintragen

```

# This is the configuration file for the 'gianduia' UR metering daemon, part of
# gLite DGAS, DataGrid Accounting System.
# This file is also used by "outOfBand scripts"
# Author: A.Guarise -- andrea.guarise@to.infn.it
# Author: R.M.Piro -- piro@to.infn.it
# Author: G.Patania -- patania@to.infn.it
#

```

```
# You should modify the following parameters according to your needs.
#
# This configuration file is read by the DGAS 'gianduia' daemon
#
# This is the directory where blapd puts the files for the LRMS jobs with
# the information needed by the accounting service
chocolateBox = "/opt/glite/var/dgasRawBox/"
#
# This is the location of the directory where gianduia puts the files
# with the usage records. It must be the same as dgasURDir specified in
# the dgas_atmClient.conf file
gianduiottiBox = "/opt/glite/var/dgasURBox/"
#
# this is the location of the directory where garbage files are stored
# for post-mortem analysis
garbageCollector = "/opt/glite/var/garbageCollector/"
#
# the lock file for glite-dgas-gianduia.pl
lockFileName = "/opt/glite/var/dgas_gianduia.lock"
#
# the log file name for glite-dgas-gianduia.pl
logFileName = "/opt/glite/var/log/dgas_gianduia.log"
#
# the lock files for glite-urcollector.pl and HAD
collectorLockFileName = "/opt/glite/var/dgas_gianduia_urCollector.lock"
collectorHADLockFileName = "/opt/glite/var/dgas_gianduia_urCollector_HAD.lock"
#
# the log file for glite-urcollector.pl
collectorLogFileName = "/opt/glite/var/log/dgas_gianduia_urCollector.log"
#
# the lock file for out of band jobs:
oobLockFileName = "/opt/glite/var/dgasOOB.lock"
#
# the buffer file that keeps track of the last processed LRMS job_id and file:
# for dgasOutOfBand_(pbs/lsf).pl
bufferFileName = "/opt/glite/var/dgasOOBuffer"
```

```

# the buffer file that keeps track of the last processed LRMS job_id and
# log timestamp for glite-urcollector.pl:
# (Note: the LRMS type specified in lrmsType will automatically be appended
# as suffix, for example: "/opt/glite/var/dgasCollectorBuffer.pbs")
collectorBufferFileName = "/opt/glite/var/dgasCollectorBuffer"

# chocolate Box parse interval (if all jobs have been processed: seconds to wait
before looking for new jobs in chocolateBox):
mainPollInterval = "60"

# if still jobs to process: number of jobs to process at each processing step
# (several steps per mainPollInterval, depending on the number of jobs found
# in chocolateBox):
jobPerTimeInterval = "10"

# time in seconds to sleep after each processing step (if there are still
# jobs to process, otherwise start a new mainPollInterval):
timeInterval = "5"

# garbage clean-up interval in seconds
queuePollInterval = "600"

#siteName parameter can be used to assign to the UsageRecord a tag identifying
# the site. Example: "INFN-TORINO".
siteName = "Example-RZ"

# This defines whether to use the name of the host on which the sensors run
# as the host name of the CE (in case of local jobs), default is "no":
localHostNameAsCEHostName = "no"

# This parameter specifies the LRMS type and _has to be set_. Possible values:
# "pbs", "lsf".
lrmsType = "pbs"

# This is the location of the directory where PBS accounting logs are stored.
pbsAcctLogDir = "/var/spool/torque/server_priv/accounting/"

```

```

# This is the location of the directory where LSF accounting logs are stored.
lsfAcctLogDir = "{path_lsf}/mnt/work/{cluster name}/logdir"

# The following option defines whether urcollector shall consider
# a) "grid": ONLY grid jobs, recognized by their presence in the CE's job map
#           log file (this REQUIRES useCEJobMap = "yes" and ceJobMapLog to be
#           set!)
# b) "local": ONLY local jobs, recognized by their absence in the CE's job map
#           log file (this also REQUIRES useCEJobMap = "yes" and ceJobMapLog
#           to be set!)
# c) "all": [default] all jobs; grid-related information (apart from the VO)
#           is determined only if useCEJobMap = "yes" and ceJobMapLog are set!
#           (and of course only for jobs that have such information in the
#           CE's log)
#jobsToProcess = "all"

# This is the location of the CE's grid user/grid job -> local job ID map.
# It has to be specified _either_ as a directory (containing file names such as
# "20050307" or "2005-03-07") _or_ as a file name prefix (including path)
#
# Examples for matching log file names:
#
# "<ceJobMapLog>20060309(.gz)" (one file per day; as prefix)
# "<ceJobMapLog>(- or _)20060309(.gz)" (one file per day; as prefix)
# "<ceJobMapLog>/20060309(.gz)" (one file per day: in dir)
#
# "<ceJobMapLog>2006-03-09(.gz)" (one file per day; as prefix)
# "<ceJobMapLog>(- or _)2006-03-09(.gz)" (one file per day; as prefix)
# "<ceJobMapLog>/2006-03-09(.gz)" (one file per day; in dir)
#
# "<ceJobMapLog>(.gz)" (single file; as prefix)
# "<ceJobMapLog>.1(.gz)" (single file rotated; as prefix)
#
# IMPORTANT NOTE: The CE's job log file will be used only if
# useCEJobMap is set to "yes" (default!), otherwise no grid-related

```

```

# information (user DN, user FQAN, grid job ID, etc) will be added
# (jobs considered to be local jobs).
# useCEJobMap = "yes"
useCEJobMap = "no"
ceJobMapLog = ""

keyList = "GlueHostBenchmarkSF00,GlueHostBenchmarkSI00"

#ldifDefaultFiles = "file1,file2"
# glueLdifFile = "/opt/glite/etc/glite-ce-ce-plugin/out.ldif"
glueLdifFile = "/opt/glite/etc/dgas_cpu_bench.conf

dgasCeServerdLog = "/opt/glite/var/log/dgasCeServerd.log"
dgasCeServerdLock = "/opt/glite/var/dgasCeServerd.lock"
dgasCeServerdHadLock = "/opt/glite/var/dgasCeServerdHad.lock"

```

## 5.5 dgas\_VOUserHLR.conf

Standardkonfiguration, es wurde nichts geändert!

```

# This list may contain default User HLR contact strings to be used for
# different VOs, if the HLRLocation is not specified in the JDL expression.
# Each line should have the following format:
# <vo_name> = "<HLR_contact_string>"
# where <vo_name> has to be _exactly_ the VO name as specified in the FQAN
# of the user's VOMS certificate and <HLR_contact_string> is a string of the
# form: "<HLR_hostname>:<port>:<HLR_host_certificate_subject>"
#
# Example:
#
#                               alice
"grid004.to.infn.it:56568:/C=IT/O=INFN/OU=Host/L=Torino/CN=grid004.to.infn.it"

```